

AS



Responsibility and Failure

Author(s): John Martin Fischer

Source: *Proceedings of the Aristotelian Society*, New Series, Vol. 86 (1985 - 1986), pp. 251-270

Published by: [Blackwell Publishing](#) on behalf of [The Aristotelian Society](#)

Stable URL: <http://www.jstor.org/stable/4545047>

Accessed: 24/02/2011 13:39

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=aristotelian>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The Aristotelian Society and Blackwell Publishing are collaborating with JSTOR to digitize, preserve and extend access to Proceedings of the Aristotelian Society.

<http://www.jstor.org>

XIV*—RESPONSIBILITY AND FAILURE

by John Martin Fischer

I wish to defend the thesis that there is a basic asymmetry between moral responsibility for *actions* and moral responsibility for *omissions*. One can be morally responsible for performing an action, even though one couldn't have avoided performing it—moral responsibility for action does not require freedom to do otherwise. But one can't be morally responsible for failing to perform an act, if one couldn't have performed it—moral responsibility for things we have left undone (omissions) requires freedom to do the thing in question. Thus, there is a fundamental difference between actions and omissions, with respect to moral responsibility. In this paper, I shall first discuss some examples which motivate the asymmetry claim and state what I believe the significance of the claim is. I shall then discuss and criticize a strategy, recently presented by Harry Frankfurt, which denies the asymmetry claim in the general, unqualified form in which I have stated it.¹ Finally, I shall defend the asymmetry claim by showing that it is derivable from an extremely plausible principle which connects moral responsibility with *control*. We believe that moral responsibility is associated with control, and it is in virtue of this fact that it can be seen that there is a basic difference between moral responsibility for acts and moral responsibility for omissions.

I

Acts and Omissions

Suppose that, unbeknownst to Jones, a device has been implanted in his brain by a team of scientists. This device (a sophisticated device which is, as yet, not on the market!) monitors his brain activity and, by employing it, the scientists can ascertain Jones's mental states—his desires, beliefs, choices, decisions, etc. The device also allows the scientists to induce in

* Meeting of the Aristotelian Society held at 5/7 Tavistock Place, London WC1, on Monday, 16 June, 1986 at 6.00 p.m.

¹ Harry Frankfurt, 'What We are Morally Responsible For', in L. Cauman (*et al.*), eds., *How Many Questions?*, (Indianapolis, Hackett, 1982), pp. 321-335.

Jones whatever mental states they want him to have by direct electronic stimulation of certain key parts of Jones's brain. If the scientists want Jones to choose to perform a certain act and then to act on this choice, they can bring this about by electronic manipulation of the brain.

Now imagine that Jones is walking along a beach, and he sees a child struggling in the water. Jones knows that he can swim well, and he believes that he can save the child with only minimal effort and inconvenience. After brief deliberation, Jones dives in and rescues the child. Also imagine that the scientists want Jones to act in this way, and that, in this case, their device simply monitors Jones's brain activity but plays no causal role in either Jones's deliberation or his action. Had Jones shown an inclination to choose *not* to save the child, the scientists would have activated the device and thus would have brought it about electronically that Jones choose to save the child (and that he do so), but as things actually work out, the device plays no role in Jones's decision or action.

I believe that Jones is morally responsible for saving the child—his act reflects something about Jones (perhaps a character trait) which is commendable, for which he can reasonably be praised. And yet Jones couldn't have done otherwise. Had he been inclined to choose not to save the child, he would have been *made* to choose to save the child (and to act on this choice). Thus, it seems that Jones is morally responsible for an act which he couldn't avoid doing. Mere inability to do otherwise does not rule out moral responsibility for an action; it might be the case that the fact that a person couldn't have done otherwise is *irrelevant* to the explanation of the action and thus that the person is morally responsible, even though he couldn't have done otherwise.²

²For such examples, see Harry Frankfurt, 'Alternative Possibilities and Moral Responsibility', *Journal of Philosophy*, December 1969; and John Martin Fischer, 'Responsibility and Control', *Journal of Philosophy*, January 1982. In this article, I used the term 'control', to mean 'freedom to do otherwise'. In 'Responsibility and Control', I have defended the claim that (for instance) Jones is morally responsible for saving the child although he couldn't have done otherwise. (In 'Ability and Responsibility', *The Philosophical Review*, 87, (1978), pp. 201-224 and *An Essay On Free Will* (Clarendon Press, Oxford, 1983), pp. 161-182, Van Inwagen argued against this claim.) Whereas it is inappropriate to go into the details of this debate here, I shall say a few words. One line of attack developed by Van Inwagen claims that, although Jones is responsible for the act-

Consider a second case. This time Jones does not have any fancy mechanism in his brain, and he is again strolling along the beach. Again, he sees a child struggling in the water, and he believes that he can rescue the child with minimal effort and inconvenience. But he decides not to go to the trouble, and he continues walking along the beach. The child drowns. Unbeknownst to Jones, there is a school of sharks patrolling the water between the beach and the struggling child, and had he jumped into the water, the sharks would have attacked and eaten him.

Is Jones responsible for failing to rescue the child? I think *not*. After all, he *couldn't* have rescued the child. And though the fact that he couldn't have rescued the child played no role in his decision and his omission, it is precisely this fact (that he couldn't have done otherwise) that makes me think that he is *not* responsible for his failure to save the child. Of course, in his behavior, Jones manifests something bad about himself, and he *is* morally responsible for *something*. He is, for instance, morally responsible for his failure to *try* to save the child (and his failure to jump into the water, etc.) But he *could* have tried to save the child (and he could have jumped into the water, etc.) This case is like one in which Jones sees a child in the water who (unbeknownst to Jones) has just died. When Jones fails to attempt to rescue this child, he may well be morally responsible for failing to try to save the child (and failing to swim to the child, etc.); but he is clearly not morally responsible for failing to save the child—the child is, after all, dead (at the pertinent time).

Here is another case. Smith is in his apartment, looking out of the window. He sees a man being mugged on the sidewalk in front of his building. He quickly proceeds to call the police and report the incident. The police come quickly, and the mugger is

token of his saving the child, he *could* have done otherwise, since his saving the child as a result of the scientist's intervention would be a *different* act-token from the actual one. But I argue (in 'Responsibility and Control') that, even if this were true, it seems to conflate the *possibility* of another event-token's taking place with *ability* to do otherwise, in the sense pertinent to responsibility. When we associate moral responsibility with freedom to do otherwise, we associate it with *deliberate* control, not the mere possibility of another event's taking place. And Jones does not have deliberate control—he cannot make another event occur as a result of an appropriate sort of intention.

arrested. (Obviously, this case does not take place in New Haven!) Unbeknownst to Smith, had he been inclined to choose not to call the police, he would have been seized by such intense feelings of guilt that he would have chosen to call the police (and would have called them) anyway. (I assume that the guilt feelings would have been irresistible in this case.) It seems to me that Smith's act of calling the police is laudable—he is morally responsible for calling the police. But he couldn't have avoided doing so; he would have been made to call anyway by his strong sense of guilt. Insofar as the fact that he couldn't have done otherwise plays no role in his behavior, it is evident that Smith can be held morally responsible for calling the police. Smith is praiseworthy on the basis of his action, though he couldn't have avoided it.

Next, consider the following example. Again, Smith is in his apartment, looking out the window. He sees a man being mugged on the sidewalk in front of his building. He knows that it would be easy to dial the police number on the telephone, and he believes that he could do so with little effort or trouble. But he considers the inconvenience (perhaps he will have to go down to the police station and be a witness, etc.), and he decides not to get involved. He fails to call the police, and the mugger escapes. Unbeknownst to him, his telephone line has been cut, and his telephone is not working. So he *couldn't* have called the police. Even had he chosen to call the police, the fact that the wire was cut would have prevented him from doing so.

In this case, is Smith morally responsible for not calling the police? Again, I think *not*.³ After all, he *couldn't* have called the police, and it is precisely this fact which exerts pressure on me to think that he is not morally responsible for failing to call the police. I think that any inclination to believe that Smith is morally responsible for failing to call the police results from an ambiguity in the notion of 'failing to call'. On the one hand, 'failing to call' can mean 'failing to dial', in which case Smith *is* morally responsible for failing to call the police. But this sense of failing to call is the same as 'failing to try to call' (and Smith *could* have tried to call and he certainly could have dialed). The

³ Peter Van Inwagen comes to a similar conclusion about a similar case he presents in: *An Essay on Free Will* (Clarendon Press, Oxford, 1983), pp. 165–166.

intended sense of 'failing to call', in the example, is 'failing to dial and successfully reach' the police. (Otherwise it wouldn't be a case in which an agent is purportedly responsible for failing to do something which he *couldn't* have done.) Of course Smith acts reprehensibly and is morally responsible for *something*. Smith is morally responsible, for instance, for failing to *try* to call the police, for failing to dial the police number, etc. But Smith is *not* responsible for not calling the police, i.e., for not successfully reaching the police. Suppose, similarly, that there had been an emergency in the town, and (unbeknownst to Smith) no one was at the police station. Surely, if this were so (but it played no role in Smith's actual behavior), it would *not* be appropriate to hold Smith responsible for not *successfully reaching the police*; it would only be reasonable to condemn Smith for not *trying* to call the police, for not dialing the number, etc.

The examples presented here illustrate a general point: a person can be morally responsible for doing something which he couldn't have failed to do, but *not* for failing to do something which he couldn't have done. Whenever there is an example in which it is purportedly true that a person is morally responsible for failing to do something which he can't do, it is always appropriate to redescribe the case as one in which the person is responsible (if he is responsible at all) for failing to do something which he *can* do. With omissions, one adjusts the specification of what the agent is responsible for until one finds something the agent could have done. But with actions, the corresponding move is not appropriate.

Having articulated the asymmetry principle, I want to comment on its significance. I believe that the thesis describes a feature of our conceptual framework—our way of describing and categorizing agents and actions. But it does *not* mark a certain sort of *substantive moral* difference (concerning *degree* of moral responsibility).⁴ Let us again consider the case of Smith described above, in which, unbeknownst to Smith, the telephone line has been cut. Now compare this with a similar case in which Smith also fails to call the police, but in which the telephone line has not been cut (and Smith could have successively reached the police). In the first case, Smith is morally responsible for failing

⁴Harry Frankfurt makes this point in: Frankfurt, 1982, pp. 238–329.

to *try* to call the police (and for failing to dial, etc.) but *not* for failing to call the police, and in the second case, he is responsible for failing to *call* the police. So *what* the agent is morally responsible *for* depends on things which are quite extraneous to the agent—for instance the state of the telephone line. This might seem to introduce an unacceptable kind of ‘moral luck’ into our system of evaluating agents. But, whereas our way of specifying what Smith is responsible for depends on such factors, a certain moral evaluation of Smith does *not*. Smith would be equally morally blameworthy in either case, and it would seem appropriate to apply the same kind of punishment (or blame) in both cases. So, whereas a certain kind of moral luck applies to the specification of the *content* of moral responsibility, it does *not* apply to the *extent* or *degree* of blameworthiness, and it does *not* apply to the evaluation of agents.

Very roughly, what seems crucial to our moral assessment of persons and our practice of praising and blaming is the person’s motivation (and his attempt to act on this motivation). Agents who have the same intention and make the same choice and are equally conscientious in attempting to act on the choice are accessible to the same *degree* of praise or blame, even if *what* they are responsible *for* is different. The examples show that the nature and extent of moral praise or blame do not vary in a straightforward way with changes in the specification of what the person is responsible for: *degree* of praiseworthiness or blameworthiness needn’t vary with *content* of moral responsibility.

There are other examples which seem to me to illustrate the same sort of point. Consider two truck drivers who are equally blameworthy in drinking too much and driving. They are equally reckless in running a stop sign. However, only one of them runs over a pedestrian (there is no pedestrian in the intersection when the other drunken truck driver drives through). Now it seems to me that both drivers are equally ‘bad’—their characters are equally blameworthy. Further, it seems to me that an ideally fair system of liability would require that they both pay equally into a fund to compensate the victim; after all, it is pure luck that one driver (and not the other) hits the pedestrian. Both drivers are equally blameworthy, and ideally, their liability ought to be equal; but clearly the *content* of their moral responsibility is different. One truck driver is

morally responsible for running over the pedestrian, whereas the other is *not*.⁵

I have claimed that factors quite extraneous to the agent can alter the *content* of moral responsibility, though *not* necessarily its degree. There is an analogy here between the content of moral responsibility and the content of belief. In Putnam's famous thought-experiment involving 'Twin Earth', I and my doppelganger are psychologically equivalent in all phenomenological respects, and yet my beliefs are about water, whereas his beliefs are about 'XYZ'. The *content* (in one sense, at least) of belief depends on factors quite extraneous to the agent in a way which is parallel to the way in which the content of responsibility depends on such factors. And yet there *is* something which is *shared* by me and my doppelganger—perhaps, 'phenomenal belief', or 'narrow content' (content relevant to explanation of behavior).⁶ This sameness between me and my doppelganger parallels the sameness in the *characters* of the two truck drivers. With both the cases of moral responsibility and belief, one operative notion is determined 'internally' and another operative notion is determined 'externally'. There is a notion of the content of moral responsibility, which, like the content of belief, 'ain't in the head'.

II

Frankfurt's Response

I believe that the examples discussed above show that there is an asymmetry between actions and omissions, with respect to the relevance of alternative possibilities to the specification of what the agent is morally responsible for, that is, to the *content* of moral responsibility. One way of denying the asymmetry claim presented here is to divide omissions (failures to act) into two classes and to claim that the asymmetry is true of only one class of omissions. The claim would be that when one focuses on only

⁵ There is a discussion of such cases involving moral luck, and an elaboration of the view I present here in: John Martin Fischer and Robert Ennis, 'Causation and Liability', *Philosophy and Public Affairs*, Winter, 1986, pp. 33-40.

⁶ For Putnam's thought-experiment, see: Hilary Putnam, *Mind, Language, and Reality* (London, Cambridge University Press, 1975), pp. 215-271, esp. pp. 223-227. For a discussion of the two types of content, see: Jerry Fodor, 'Methodological Solipsism considered as a Research Strategy in Cognitive Psychology', in *Representations* (MIT Press, Cambridge, Mass., 1981), pp. 225-253.

one class of omissions and then concludes that it is generally true that acts and omissions are asymmetric with respect to the requirement of alternative possibilities for moral responsibility, one is generalizing from an inappropriate sample. Let us turn to one version of this strategy, which has been presented recently by Harry Frankfurt.⁷

Frankfurt first considers the case of Smith, in which (unbeknownst to him) the telephone wire has been cut. In this case, Frankfurt agrees that Smith is not morally responsible for failing to call the police (i.e., failing successfully to reach the police). According to Frankfurt, this is because Smith's failure is an 'impersonal failure'—the failure would have occurred no matter what bodily movements Smith made.⁸ That is, no matter what bodily movements Smith would have made, they would not have had the consequence that Smith successfully reaches the police by telephone. When a failure is impersonal, Frankfurt claims that an agent is not appropriately held to be morally responsible for it.

But there is another class of failures which are 'personal'—it is *not* the case that they would occur, no matter what bodily movements the agent makes. When a failure is personal, Frankfurt claims that an agent *can* be morally responsible for failing, even though he lacks the power to perform the act in question. Frankfurt says:

. . . suppose that as Q is driving he fails to keep his eyes straight ahead because he prefers to examine the interesting scenery to his left; and suppose further that if the scenery had not distracted him something else would have brought it about that he was looking to his left at that time. In these circumstances, Q cannot keep his eyes straight ahead. Is he morally responsible for failing to do so? Of course he is! The fact that he cannot avoid failing has no bearing upon his moral responsibility for the failure, since it plays no role in leading him to fail.

Failing to keep one's eyes straight ahead is exclusively a matter of what movements a person makes; it is *constituted*

⁷ Frankfurt, 1982.

⁸ Frankfurt, 1982, pp. 330-331.

by what the person himself does, and what the person does is therefore both a sufficient and a necessary condition for it. It cannot be said, then, that Q's failure would have occurred no matter what he had done—i.e., regardless of what bodily movements he made. If he had not moved his eyes to the left at all he would not have failed.⁹

Frankfurt's point is that there is not the same reason for denying that Q is morally responsible for his failure as there is for denying that Smith is responsible for failing to call the police. Frankfurt's position, then, is that the fully general asymmetry thesis is false. One might think it is true by failing to distinguish between impersonal and personal failures and paying undue attention to impersonal failures. With respect to actions and *personal* failures, there is symmetry. According to Frankfurt, then, 'there is no inherent difference between performances and failures in virtue of which [it might be true that moral responsibility for failure to act requires the ability to act, whereas moral responsibility for action does not require the ability to fail to act]'¹⁰

I believe that Frankfurt's strategy fails. Consider again the case of Smith. It is Frankfurt's view that since Smith's failure is impersonal, he is not morally responsible for failing to call the police. But imagine that the telephone wires have not been cut; imagine, instead, that Smith has (unbeknownst to him) acquired a temporary bout of paralysis which would have made it impossible for him to walk to the phone and dial, even if he had chosen to do so. In this case, Smith's failure to call the police is *personal*: it is not the case that the failure would have occurred, no matter what bodily movements Smith made—if Smith had walked to the phone and dialed it, he would have successfully reached the police. So in the case in which Smith is paralyzed, his failure is personal, and he *can* be held morally responsible for failing to call the police (on Frankfurt's approach). But I claim that it is intuitively unacceptable to say that, in the case in which the telephone line is cut, Smith is *not* morally responsible for failing to call the police, but in the case in which he is paralyzed, Smith *can* be held responsible for failing to call the police. It

⁹ *Ibid.*

¹⁰ *Ibid.* p. 332.

seems to me that these cases should be treated in the same way; surely, the fact that, in the first case the failure is impersonal and in the second case the failure is personal, makes no difference to how we assess Smith's responsibility, and how we specify what he is responsible for. Or so it seems to me.

Further, consider again the case of Jones, who fails to save the struggling swimmer. Unbeknownst to him, there is a school of hungry sharks lurking in the water. Jones's failure to save the child is *personal*—it is not true that it would have occurred, no matter how Jones moved his body: had Jones swum to the child, he would have saved the child. It's just that, given the presence of the sharks, he *couldn't* have swum to the child. It seems that Jones is *not* responsible for failing to save the child (and for failing to swim to the child), even though his failure is *personal*. And if there were no sharks, but rather, the child had died just before Jones reached it, Jones's failure to save the child would have been *impersonal*—he would have failed to save the child, no matter what bodily movements he made. But the two cases should be treated alike. The fact that, in the first case the failure is personal whereas in the second the failure is impersonal, makes no difference to how we specify what he is responsible for.

Suppose that Frankfurt, after considering the case of Jones and the sharks and agreeing with me that Jones is not responsible for failing to save the child, modified his account of responsibility as follows: a failure is impersonal if and only if it would occur, no matter how the agent moves his body, among all the ways in which the agent *can* move his body. Now Jones's failure to save the child is indeed impersonal. But, unfortunately, on the revised account, Q's failure to keep his eyes straight ahead is *also* impersonal—it would have occurred, no matter how Q's body moved, among all the ways in which Q *could* have moved his body. Thus, the revision is unacceptable. Frankfurt's strategy of response to the fully general asymmetry thesis does not work. We have (thus far) been given no reason to think that the thesis is to be rejected.

III

Control

I have presented the asymmetry principle and criticized a strategy of response to the principle. I have simply relied on my

‘considered intuitions’ about various examples, but I have not yet developed a theoretical explanation or justification of the asymmetry claim. In this section, I shall attempt to do this, and in order to do so, I shall first consider the concept of control. It will be in virtue of the association of responsibility with control that I shall explain and defend the asymmetry thesis.

We sometimes think that a person has control over some event or state of affairs.¹¹ But there are, I believe, at least two importantly different kinds of control.¹² To illustrate these different sorts of control, let us consider a pilot who is flying a plane; he is, we might say, ‘at the controls’ of the plane. Now suppose that the pilot turns the plane’s steering apparatus to the left, because he has decided that the plane should go westward, thus causing the plane to turn westward. In this case, there is an actual causal sequence initiated by the pilot which issues in the plane’s turning westward. And let us say that the pilot deliberated prior to reaching his decision, and that his decision was based, in an ‘appropriate’ way (the way normally involved in rational decision-making), on his reasons, and that his decision caused (again, in an ‘appropriate way’) his turning the steering apparatus to the left. (The idea here is that only *certain* sorts of causal chains issuing in the event of the plane’s turning westward would be consistent with the pilot’s having control over this event.) In short, the pilot causes, in a suitable way, the plane’s turning to the west. When the pilot causes the plane to turn to the west in a suitable way, we might say that he *controls* the plane’s movement toward the west, that he has a certain sort of ‘control over’ the plane’s movement, and that he is ‘in control of’ the plane’s movement. I shall call this kind of control over an event, ‘actual causal control’ over an event; when the pilot steers the plane westward, he has actual causal control of the plane’s movement toward the west.

Now imagine that the plane is flying on ‘automatic pilot’; a

¹¹ Of course, it is not only persons who can have control; a thermostat may control the temperature of a room, for instance. And it is not only events or states of affairs over which there can be control.

¹² I am deeply indebted to Michael Bratman for suggesting the basic idea of this section to me and for helping me to develop the idea. Also, there is a similar distinction between two sorts of ‘guidance’ in: Harry Frankfurt, ‘The Problem of Action’, *American Philosophical Quarterly*, April 1978, pp. 157–162.

computerized navigational system is operating the controls of the plane, and the pilot is simply standing by and 'monitoring' the navigational system. Suppose that the computer causes the 'automatic pilot' to turn the steering apparatus to the left, thus causing the plane to go to the west—the pilot does not intervene, and plays no role in the causal sequence which results in the plane's turning westward.

There is a sense, however, in which the pilot might have control over the plane's turning westward. Imagine, first, that although the pilot doesn't actually intervene, he has the power to *ensure* that the plane turns to the west. That is, had the computer issued instructions to veer to the east (or perhaps, had the automatic pilot malfunctioned and begun to turn the plane eastward), the pilot had the power to cause (in a suitable way) the plane nevertheless to turn westward. Or had there been a strong gust of wind which began to turn the plane eastward, the pilot could have turned the steering apparatus so as to cause (in an appropriate way) the plane to turn westward. Thus, whereas the pilot plays no role in the actual causal sequence, he nevertheless has the power to react to a certain range of contingencies in such a manner as to ensure the desired outcome. The outcome is here rendered *resilient* by the pilot—it will remain the same, even if there are any number of perturbations.

Imagine also that the pilot has the power to *prevent* the plane's turning westward. That is, the pilot has the power to 'deactivate' or 'override' the automatic pilot and take over the steering apparatus himself, and he can cause the plane to turn eastward (or continue straight ahead), if he wishes. Though he actually allows the plane to turn to the west, he can prevent it. The outcome is here *responsive* to the pilot—he can change it, if he so chooses.

Now even when a person does not actually causally control an event, but has the power both to *ensure* the occurrence of the event (in the sense discussed above) and to *prevent* the occurrence of the event, then he has a certain sort of control over the event. I shall call this kind of control over an event, 'regulative control'; in the second 'pilot case', the pilot has regulative control of the plane's turning westward.

It is possible that the pilot be able to ensure the plane's

turning westward without being able to prevent it. This would be the case if the plane's steering apparatus would function properly, if it were turned to the left (the way which causes the plane to go westward), but would not function properly, if it were turned any other way—if it were turned at all, it would cause the plane to go westward. Also, it is possible that the pilot be able to prevent the plane's turning westward, without being able to ensure it. This would be the case if the plane's steering apparatus would malfunction, if turned by the pilot (rather than the automatic pilot) any way other than the way which causes the plane to fail to go westward—if the pilot were to try to cause the plane to go westward, he wouldn't do so, though he could prevent it's going westward. I think that it is clear that a person must have *both* capacities—the capacity to ensure and the capacity to prevent—in order to have regulative control; if the pilot lacked (for instance) the ability to ensure that the plane turn westward, then he wouldn't truly have *control* over it's going westward (in the 'regulative' sense). When one has regulative control over an upshot, one renders the upshot both resilient and responsive.

It is obvious, from the second case of the pilot, that an individual can have regulative control over an upshot without actually causally controlling the upshot. Also, it is possible that an agent have actual causal control of an event without having regulative control of it—the cases discussed above illustrate this. (When Jones saves the child, though he couldn't have done otherwise, he actually causally controls the child's being saved, but he lacks regulative control over the child's being saved, since he couldn't have prevented it. The situation is the same, when Smith actually calls the police, though he couldn't have done otherwise.) Finally, it is obvious that one can have (or lack) *both* kinds of control, and indeed, in many typical cases in which one is said to have control, one possess both kinds of control.

I have distinguished two sorts of control over events (or states of affairs). But what is the relationship between events (and states of affairs), on the one hand, and actions and omissions, on the other? First, I take it that it is reasonable to think that when a person performs an action—does something—he causes (or brings about) an event. So when the pilot performs the action of turning the plane to the west, he causes (in an appropriate way)

a certain upshot: the plane's turning to the west. And when an individual turns on the light, he causes (in a certain way) the light's going on, etc. In performing an action, an agent stands in the 'bringing about' relation to an event.

It might be thought that in omitting to perform an action—in failing to do something—an agent causes (or brings about) a different sort of event, a 'negative event' or perhaps, a negative state of affairs. So, in failing to turn the plane westward, the pilot might be thought to cause the plane's not turning westward. But this needn't be so; an agent might refrain from doing something while (for example) *someone else* does it, and thus, in refraining from performing the act, the agent *doesn't* bring about the pertinent 'negative' event or state of affairs. The pilot might refrain from turning the plane westward because he sees that his co-pilot is doing so (or perhaps that the automatic pilot is doing so); though the pilot fails to turn the plane westward, he does *not* cause the plane's not turning westward; after all, the plane *does* turn westward. So, with acts and omissions, we do not have *one* relation—the causing or bringing about relation—and *two* kinds of events or states of affairs—positive and negative; rather, we have *two* kinds of relations—bringing about (causing) and not bringing about (not causing)—and *one* kind of event or state of affairs. When an agent performs an action, he stands in the 'bringing about' relation to a certain upshot, and when an agent refrains from performing (the same) action, he stands in the 'not bringing about' relation to the same upshot.

I wish to explore this point a bit more carefully and clarify my conception of omissions or failures to act. There are wider and narrower conceptions of omissions. I wish to adopt a quite wide conception, according to which it is true that whenever a person doesn't do X, then he fails (in the relevant sense) to do it, and he omits to do it. Thus, I am now failing to stop the Earth's rotation (and omitting to do so). Omission to do X needn't (in this sense) require explicit deliberation about X or the ability to do X, etc. (Some philosophers have claimed that one omits to perform an act only if one can perform it—the asymmetry thesis would follow immediately from this rather strong assumption, given the claim about responsibility for *action*.) It is desirable to have a theory of responsibility which applies to this broad notion of omission, as we often speak of an agent's being (or not being)

morally responsible for not doing something (where his failure might not count as an omission, in a narrower sense of 'omission').

Now consider a case in which I refrain from shooting a baby. I've argued that the omission is to be analyzed as my standing in the not-causing relation to a certain event—the baby's being hit by a bullet. This is because I could refrain from shooting the baby, while someone else shoots it; thus, my omission can't be understood as my standing in the causing relation to 'the baby's not being hit by a bullet'. But it might be objected that I have simply chosen the 'wrong' event for the analysis; when I omit to shoot the baby, it might be claimed that I stand in the causing relation to 'the baby's not being hit by a bullet from my gun'.

But this sort of analysis will not work suitably for all omissions, that is, for all cases in which I fail to perform a given act. Suppose, for example, that I sneeze or am suddenly hit by a strong gust of wind, and this causes the event of the baby's being hit by a bullet from my gun. Now it seems clear that, whereas the baby is hit by a bullet from my gun, I do *not* shoot the baby. I do not perform any action at all. I do not shoot the baby, but this is not a case in which I cause the baby's not being hit by a bullet from my gun; after all, the baby *is* hit by a bullet from my gun. Thus, we can't analyze omissions (in the sense of 'omission' relevant here, in which it follows from an agent's not doing X that he omits to do X) by construing omission in the way suggested, i.e., as involving the 'causing' relation to certain 'finely specified' negative events. In any case, rather than positing special 'negative' events, about which there may be considerable unclarity, it is preferable to understand omissions as involving the 'not-causing' relation to ordinary events. I shall leave unspecified here what exactly these events are; for some omissions (omissions involving 'non-purely causal verbs'), the relevant event might be a more finely construed event. For instance insofar as 'shooting' is not a purely causal verb, then my shooting the fat man can't be analyzed as my causing the fat man to be shot; after all, I could cause him to be shot by paying *you* to shoot him. So, when I refrain from shooting a child, it might be appropriate to say that I stand in the 'not-causing' relation to the event, 'the child's being shot *by me*'. In general, purely causal verbs will require events which are less finely

specified. (If 'saving' is a purely causal verb, then when I save a child, I stand in the 'causing' relation to the event, 'the child's being saved', and when I refrain from saving a child, I stand in the 'not-causing' relation to the event, 'the child's being saved'.) Of course even when the 'finer' event is used, there is a crucial difference between construing my refraining from shooting the baby as my causing the baby's not being hit by a bullet from my gun and my not causing the baby's being hit by a bullet from my gun.¹³

We now have the two ingredients required for the defense of the asymmetry principle: (i) the distinction between two kinds of control, and (ii) the analysis of the nature of the relations involved in action and omission. Given the claim that moral responsibility for *action* does not require freedom to do otherwise, the basic principle from which the asymmetry thesis can be derived is: moral responsibility requires *control* (of at least *one* of the two sorts).

Let me elaborate. We begin by considering a typical case where an agent can be said to be morally responsible for performing an action. Let us say that, under ordinary circumstances (no manipulative scientists, sharks, etc.) a lifeguard jumps into the water and saves a drowning child. The lifeguard is rationally accessible to praise for saving the child—he is morally responsible for saving the child. The lifeguard, insofar as he saves the child, stands in the 'bringing about' relation to the upshot, 'the child's being saved'. And in this case, the lifeguard has (on plausible assumptions) both actual causal control and regulative control over the relevant event—the child's being saved. Obviously, other cases could be presented in which an agent is morally responsible for performing an act and the agent has both actual causal control and regulative control over the relevant event.

What is interesting about the kinds of examples discussed above—'Frankfurt-style' counterexamples to the principle that

¹³ In the analysis I have sketched, 'not causing' is to be understood as 'not-causing-in-the-appropriate-way'. So, if a gust of wind causes me to pull the trigger, I do *not* cause in the way appropriate to control the death of the baby. I do not purport here to give an entirely adequate analysis of causal verbs (and of omissions); this is beyond the scope of this paper. My main point is that it is more promising to construe omission as involving the not-causing relation to ordinary events than the causing relation to negative events.

moral responsibility for an act requires freedom to do otherwise—is that they show that an agent can be morally responsible for performing some act X even though he lacks regulative control over the pertinent event, X'. It seems that in these examples the agent is responsible for doing X in virtue of actually causally controlling X'. That is, moral responsibility for doing X requires control of *some* sort over X'. Above, I claimed that doing X and failing to do X both involve relations (of different kinds) to the same (possible) event, X'. So it is extremely plausible to suggest the following basic principle of moral responsibility: moral responsibility for doing X or for failing to do X requires at least one of the two kinds of control over X'. At a deep level, this principle treats actions and omissions uniformly—responsibility for doing X or for failing to do X both require control over X'. But the asymmetry thesis can be easily derived from this basic principle (together with the intuitions about responsibility for actions presented above), as I shall now show.

According to the basic principle, moral responsibility requires at least one of the two kinds of control. And there are cases in which an agent is morally responsible for doing X, though he lacks regulative control over X'. But by the very nature of omissions, when an agent omits to perform X, he does *not* actually causally control the relevant event, X'. That is, if the lifeguard fails to save the child, he does not actually causally control the child's being saved—he does not cause (in the appropriate way) the child to be saved. Thus, when an agent omits to perform X, he *must* have *regulative* control over X', if he is to have *any* sort of control over X'. So if the lifeguard fails to save the child, he must have regulative control over the child's being saved, if he is to have any sort of control over the child's being saved. In general, because of the very nature of omissions, when an agent omits to do X, he doesn't actually causally control X'; and thus, he must have regulative control over X', if he is to have any sort of control over X'. And thus, by the basic principle of moral responsibility, if an agent is to be morally responsible for omitting to do X, he must have regulative control over X'.

If the lifeguard is to be morally responsible for failing to save the child, he must have regulative control over the child's being

saved. But in order to have regulative control over the child's being saved, he must have the power to *ensure* that the child is saved, i.e., to cause (in the manner appropriate to control) the child's being saved. Thus, he must be able to save the child. In general, then, when an agent omits to do X, he must be able to do X, if he is to have any sort of control over X'. And so, in general, if an agent is to be morally responsible for failing to do X, he must be able to do X. But when an agent does X, he may exhibit direct causal control without having regulative control over the upshot. Thus, the asymmetry principle can be derived from the basic principle of moral responsibility.

IV

Conclusion

I have used two ingredients to generate a basic principle concerning moral responsibility: (i) the claim that action and failure to act involve different sorts of relations (causing or not-causing) to the same sort of event, and (ii) the analysis of the two types of control—actual causal control and regulative control. The basic principle states that moral responsibility is associated with control: to be morally responsible for either doing or failing to do X, one must have *some* kind of control over X'. Given this basic principle and the view that responsibility for *action* doesn't require the freedom to fail to act, we can derive the asymmetry claim. Whereas Susan Wolf has recently argued that there is an asymmetry with respect to good and bad acts, I have argued that there is an asymmetry with respect to acts and omissions.¹⁴

I wish to make two points. First, I reiterate the point that the asymmetry thesis does *not* entail that the *degree* of praiseworthiness or blameworthiness of an agent who fails to perform an act depends on the existence of alternative possibilities; my point concerns the *content* of responsibility rather than its degree, and it does not pertain to the assessment of *agents*.

Second, I have obviously not 'proved' the basic principle of moral responsibility, from which I have derived the asymmetry principle. Indeed, I don't see how one would go about proving it. Rather, I have suggested it as a plausible and attractive principle which embodies, at a deep level, a *symmetric* account of

¹⁴Susan Wolf, 'Asymmetrical Freedom', *Journal of Philosophy*, March, 1980.

moral responsibility. The asymmetry thesis can thus be seen to derive from a deeper symmetry, and to this extent, is not *ad hoc* and mysterious. I have attempted to develop a way of illuminating what (to me, at any rate) are persistent and resilient intuitions of asymmetry—intuitions which seem, at first, bizarre and inexplicable. I have, then, undertaken to answer Frankfurt's challenge to specify the 'inherent difference' between actions and omissions in virtue of which moral responsibility for failures to act requires the ability to act, whereas moral responsibility for actions does not require the ability to fail to act.

I wish to end with a few remarks about causal determinism and moral responsibility. It might be thought that Frankfurt-type examples in which an agent is morally responsible for an action although he couldn't have done otherwise show that causal determinism is compatible with moral responsibility for action, even if determinism is incompatible with freedom to do otherwise. In 'Responsibility and Control', I argue that even if there are cases in which one is morally responsible for an act which one couldn't have avoided performing, this doesn't yet establish compatibilism about determinism and responsibility for action. Frankfurt-style examples should incline us to accept an 'actual-sequence' approach to responsibility; there seem to be cases in which the fact that the agent couldn't have done otherwise plays *no role* in the actual sequence issuing in action (though it *would have* played a role in some alternative sequence), and thus, in which the agent is morally responsible for the action although he couldn't have done otherwise. But causal determinism is a doctrine about what occurs in the actual sequence, and thus it might be argued that if causal determinism obtains, then the fact that the agent couldn't have done otherwise *does* play a role in the sequence actually issuing in action. So even if there are (or might be) cases in which a person is morally responsible for performing an action which he couldn't have avoided performing, this wouldn't in itself show that there are (or might be) cases in which a person is morally responsible for performing an action which is causally determined. Thus the examples don't rescue our intuitive, commonsense notion of responsibility from the threat posed by determinism.

In this paper I have articulated another reason why

Frankfurt-style examples concerning actions don't rescue our commonsense framework of responsibility. *Even if* these examples are cases in which a person is morally responsible for performing an action which he couldn't have avoided performing, there don't seem to be any such cases in which an agent is morally responsible for failing to perform an action which he couldn't perform. But on our ordinary notion of responsibility, we can be morally responsible for *failures* as well as actions. So if determinism threatens our ability to *do* things we have left undone, then determinism still threatens our ordinary notion of responsibility.¹⁵

¹⁵ I am grateful for comments by Carl Ginet. Also, I benefitted greatly from reading an earlier version of this paper at the University of California, Santa Barbara. I am particularly indebted to a number of useful discussions with Michael Bratman. My work on this paper has been supported by a Fellowship for Independent Study and Research from the National Endowment for the Humanities (U.S.A.).